Beyond histograms: efficiently estimating radial distribution functions via spectral Monte Carlo

Thomas W. Rosch and Paul N. Patrone*
National Institute of Standards and Technology
(Dated: September 5, 2016)

Despite more than 40 years of research in condensed-matter physics, *state-of-the-art* approaches for simulating the radial distribution function (RDF) $g(r)$ still rely on binning pair-separations into a histogram. Such methods suffer from undesirable properties, including subjectivity, high uncertainty, and slow rates of convergence. Moreover, such problems go undetected by the metrics often used to assess RDFs. To address these issues, we propose (I) a spectral Monte Carlo (SMC) method that yields $g(r)$ as an analytical series expansion; and (II) a Sobolev norm that assesses the quality of RDFs by quantifying their fluctuations. Using the latter, we show that, relative to histogram-based approaches, SMC reduces by orders of magnitude both the noise in $g(r)$ and the number of pair separations needed for acceptable convergence. Moreover, SMC reduces subjectivity and yields simple, differentiable formulas for the RDF, which are useful for tasks such as coarse-grained force-field calibration via iterative Boltzmann inversion.

PACS numbers: 61.20.Ja
Keywords: Radial distribution function, Monte Carlo, molecular dynamics, iterative Boltzmann inversion

arXiv:1607.03429v2 [cond-mat.mtrl-sci] 2 Sep 2016

In simulations of condensed matter systems, one can barely overstate the importance of the radial distribution function (RDF) $g(r)$. To name only a few applications, $g(r)$ is used to (i) link thermodynamic properties to microscopic details [1–3]; (ii) compute structure factors for comparison with X-ray diffraction [4, 5]; and more recently, (iii) calibrate interparticle forces for coarse-grained (CG) molecular dynamics (MD) [6–11]. Indeed, the RDF is such a key property that in the past few years, much work has been devoted to estimating $g(r)$ via parallel processing on GPUs [12]. Given these observations, it is thus surprising that state-of-the-art techniques still construct $g(r)$ by binning simulated pair-separations into histograms, with little thought given to developing more efficient methods [3, 13].

In this letter, we address this issue by proposing a spectral Monte Carlo (SMC) method for computing simulated RDFs. The key idea behind our approach is to express $g(r)$ in an appropriate basis set and determine the mode coefficients via Monte Carlo estimates. Relative to binning, we show that this approach decreases subjectivity of the analysis, thereby reducing both the noise in $g(r)$ and the number of pair separations needed to generate useful RDFs. To support these claims, we also discuss how traditional L^2 (or sum-of-squares) metrics are insufficient for assessing convergence of $g(r)$ and propose a Sobolev norm [14] as an appropriate alternative.

The motivation for this work stems from the fact that $g(r)$ is increasingly being used in settings in which the details of its functional form play a critical role. For example, scientists now routinely simulate untested materials in an effort to tailor their structural properties without the need for expensive experiments [15, 16]; in such applications, *objectively* computing RDFs is a key task. Along related lines, structural properties are increasingly being used to calibrate coarse-grained force-fields [6–11].[1] The success of such strategies often relies on being able to differentiate $g(r)$, which requires that simulated RDFs be accurate and relatively noise-free.

In this light, we therefore emphasize that histogram-based RDFs suffer from an inability to objectively control uncertainties. This arises for several reasons. For one, histogram bin-sizes are *subjective* parameters that limit the resolution of small-scale features, and often one must trade this resolution for reduced noise. Smoothing is sometimes used as an alternative to increasing bin-sizes, but this introduces difficult-to-quantify uncertainties that depend on the choice of method. Moreover, finite differences and/or derivatives are known to amplify noise, which renders tasks such as CG force-field calibration more difficult. Given that (i) a corresponding experimental RDF may be unavailable for comparison, and (ii) simulation resources are often at a premium, histogram-based approaches therefore place undue burden on modelers to obtain accurate results.

These observations therefore motivate us to propose

$$g(r) \approx g_M(r) = \sum_{j=0}^{M} a_j \phi_j(r) \tag{1}$$

where $\phi_j(r)$ are orthogonal basis functions on the domain $[0, r_c]$, r_c is a cutoff radius beyond which we do not model $g(r)$, a_j are coefficients to-be-determined, and M is a

* paul.patrone@nist.gov

[1] In iterative Boltzmann inversion, this is achieved by updating the ith correction to the CG forces $F(r)$ and energies $U(r)$ via $U_{i+1}(r) = U_i(r) + k_B T \ln[g_i(r)/g_t(r)]$, $F_i(r) = -\nabla U_i(r)$, and $g_i(r) = g_i(r, \mathbb{S}[F_i])$, where $k_B T$ is the temperature, $U_0(r) = -k_B T \ln[g_t(r)]$ for a target RDF g_t, and $g_i(r)$ is computed from a CG MD simulation $\mathbb{S}$ that uses $F_i(r)$ as the CG force [6–11].

mode cutoff. Formally, the a_j are given by

$$a_j = \int_0^{r_c} \mathrm{d}r\ \phi_j(r) g(r) = \int_0^{r_c} \mathrm{d}r\ \phi_j(r) \frac{N(r)}{4\pi r^2 \rho}, \tag{2}$$

where ρ is the bulk number density and $N(r)\mathrm{d}r$ is the expected number of particles in a spherical shell with radius r, thickness $\mathrm{d}r$, and a particle at the origin. In practice, Eq. (2) cannot be evaluated analytically, since $N(r)$ is unknown. However, MD simulations yield random pair-separations distributed according to $N(r)\mathrm{d}r$. Thus, we replace Eq. (2) by its Monte Carlo estimate [17]

$$a_j \approx \bar{a}_j = \frac{\mathcal{N}(r_c)}{n_{\mathrm{pairs}}} \sum_{k=1}^{n_{\mathrm{pairs}}} \frac{\phi_j(r_k)}{4\pi r_k^2 \rho}, \tag{3}$$

where $\mathcal{N}(r_c)$ is the expected number of particles in a sphere of radius r_c (given a particle at the origin), r_k is the kth pair separation, and n_{pairs} is the total number of such separations.

In order to simplify Eq. (3), note that $n_{\mathrm{pairs}} = n_c n_{\mathrm{ppc}}$, where n_c is the number of MD configurations (i.e. timesteps or "snapshots") used to compute $g(r)$, and n_{ppc} is the number of pairs-per-configuration. The latter is well approximated by

$$n_{\mathrm{ppc}} \approx \mathcal{N}(r_c)\mathcal{N}_{\mathrm{tot}}/2, \tag{4}$$

when $\mathcal{N}_{\mathrm{tot}}$ (the number of particles per configuration) and r_c are large.[2] Substituting Eq. (4) into Eq. (3) yields[3]

$$\bar{a}_j = \frac{2}{\mathcal{N}_{\mathrm{tot}} n_c} \sum_{k=1}^{n_{\mathrm{pairs}}} \frac{\phi_j(r_k)}{4\pi r_k^2 \rho}. \tag{5}$$

We emphasize that, as opposed to histogram-based approaches, Eq. (5) provides more objective control over uncertainties in simulated RDFs. Specifically, for many choices of $\phi_j(r)$, the mode coefficients decay as $|a_j| < Cj^{-p}$, where the constant C and rate p depend on the smoothness of $g(r)$. Furthermore, for such bases, $g_M(r)$ converges to $g(r)$ uniformly in M [19].[4] This implies that in principle, the maximum error in $g_M(r)$ is controlled through M. However, Monte Carlo sampling also introduces uncertainty in a_j, which can be estimated via

$$\sigma_j^2 = \frac{4}{(\mathcal{N}_{\mathrm{tot}} n_c)^2} \sum_k \left[\bar{a}_j - \phi_j(r_k)/4\pi r_k^2 \rho\right]^2. \tag{6}$$

[2] This identity arises as follows. First, the total number of pair separations is $\binom{\mathcal{N}_{\mathrm{tot}}}{2} \approx \mathcal{N}_{\mathrm{tot}}^2/2$ when $\mathcal{N}_{\mathrm{tot}} \to \infty$. Only considering pairs separated by $r \le r_c$, we reduce the total number of pairs by a factor of $\mathcal{N}(r_c)/\mathcal{N}_{\mathrm{tot}}$. We require r_c to be large enough so that the relative fluctuations in $\mathcal{N}(r_c)$ are small.

[3] Interestingly, related methods have been developed for density-of-state calculations under the name "kernel polynomial method." See, e.g. Ref. [18].

[4] Uniform convergence of $g_M(r)$ to $g(r)$ means that for any ϵ, there is an M such that $|g_M(r) - g(r)| < \epsilon$ holds for all r. Moreover, if $g(r)$ has p derivatives, often $|a_j| \le \mathcal{O}(j^{-p})$; if $g(r)$ has infinitely many derivatives, the $|a_j|$ usually decay exponentially.

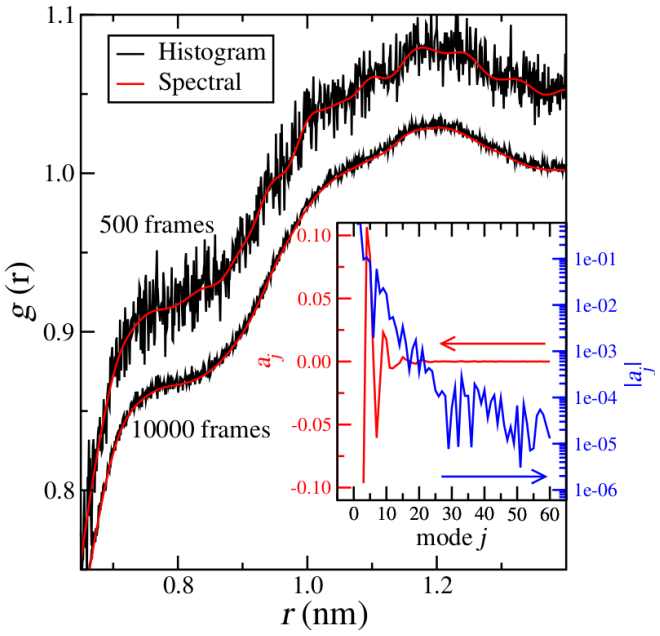

FIG. 1. RDF of atomistic polystyrene (PS) in CG coordinates using the histogram method (black, rough curves) and SMC (red, smooth curves). The upper and lower pairs are calculated with $n_c = 500$ (shifted up by 0.05) and $n_c = 10^4$ snapshots. The inset displays the spectral coefficients a_j (left scale) and $\log|a_j|$ (right scale) for the first 60 modes.

This suggests that the largest meaningful mode cutoff $M^\star$ can be estimated from $|a_M^\star| = \mathcal{O}(\sigma_M^\star)$, which corresponds to the noise-floor of $\bar{a}_j$; cf. Fig. 1. Given the uniform convergence of Eq. (1), we then conclude that: (i) the error in $g_M(r)$ is the greater of either $\mathcal{O}(\sigma_M)$ or $\mathcal{O}(a_M)$ for any cutoff; and (ii) $g_M(r)$ can model all features whose characteristic size is greater than r_c/M.

We also emphasize that the task of choosing a suitable basis is generally straightforward. It is well known, for example, that if $g(r)$ is twice differentiable and $g'(0) = g'(r_c) = 0$ (which should approximately hold if r_c is large enough), then $\phi_j(r) = \sqrt{2/r_c}\cos(j\pi r/r_c)$ converges uniformly and yields a series whose derivative converges to $g'(r)$ [20]. Moreover, $a_j \le \mathcal{O}(j^{-2})$, although exponential convergence is expected when $g(r)$ is infinitely differentiable (cf. Fig. 1) [19]. Orthogonal polynomial bases (e.g. Legendre or Chebyshev) are also reasonable, as they provide uniform approximations and similar rates of convergence, irrespective of boundary conditions [19].

In order to illustrate the usefulness of Eq. (5), we compute the RDFs of a CG polystyrene (PS) model. We first run a 10 ns, atomistic NVT simulation of amorphous, atactic PS (10 chains of 50 monomers) interacting through the pcff forcefield [21] at 800 K and $\rho = 0.758$, with configurations output every 1 ps. This trajectory is then mapped into CG coordinates at a resolution of 1 CG bead per monomer (located at the center of mass), so that $\mathcal{N}_{\mathrm{tot}} = 500$. Next, we calculate CG RDFs via a histogram with 1400 bins on the interval $0 \le r \le 1.4$ nm

and SMC with a cosine basis.

Figure 1 shows the results of these computations for $n_c = 500$ and $n_c = 10^4$. The benefits of the spectral approach are readily apparent, especially when $n_c = 500$. For $n_c = 10^4$, noise in the histogram method decreases by roughly a factor of 4 or 5 (as expected from the central limit theorem), but SMC is still dramatically smoother. The inset displays the first 60 spectral coefficients when $n_c = 10^4$. By eye, only 35 to 50 modes are required to reach the noise floor, far fewer degrees of freedom than the 1400 histogram bins.

To further illustrate the smoothness of $g_M(r)$, we use iterative Boltzmann inversion (IBI; cf. Footnote 1) to calibrate CG forces for PS using first the histogram method and then SMC. For the latter, we took $M = 60$ and computed all forces F_i analytically. For the histogram reconstruction, we used a central finite-difference scheme to approximate the F_i. IBI updates were performed using the $n_c = 10^4$ RDFs in Fig. 1 as the target RDF g_t (cf. Footnote 1). Figure 2 shows the results of these computations. Notably, the top subplot shows that after five iterations of IBI, the histogram-based force has extreme, high-frequency noise (despite taking $n_c = 10^4$), whereas the SMC force does not.

To make this comparison more quantitative, we define

$$||g||^2_{L^2} = \frac{1}{r_c}\int_0^{r_c} \mathrm{d}r\; g(r)^2 \approx \sum_{j=1}^{n_{\mathrm{bins}}} \frac{g_j^2 \Delta r_j}{r_c} \tag{7}$$

where the sum is used for the histogram reconstructions, g_j is the RDF evaluated in the jth bin, n_{bins} is the number of bins, and Δr_j is the width of the jth bin. Many works invoke $||g - g_t||_{L^2}$ (or variants thereof) to assess when a given RDF is sufficiently converged to g_t [6, 22, 23]. However, Fig. 2 shows that both the histogram and SMC RDFs converge in L^2 to their respective g_t at about the same rate, suggesting that this norm is not strongly affected by high-frequency fluctuations. To account for such effects, we propose a Sobolev norm [14]

$$||g||^2_{H^1} = ||g||^2_{L^2} + ||g'(r)||^2_{L^2}, \tag{8}$$

where we approximate $g'(r) \approx (g_{j+1} - g_{j-1})/(r_{j+1} - r_{j-1})$ for the histogram reconstructions (r_j are the bin centers). Physically, the second term of Eq. (8) assesses how smoothly $g \to g_t$. This extra information reveals a stark difference between the histogram and SMC reconstructions insofar as the former does not improve in an H^1 sense. Moreover, given that the H^1 and L^2 norms of the SMC reconstruction quickly overlap, it is clear that the difference with the H^1 norm of the histogram reconstruction is due to its high-frequency content.

To test the robustness of SMC and compare with smoothing techniques, we also used 120 cosine modes to construct the O-O $g(r)$ for an TIP4P water [24]; cf. Fig. 3. The atomistic system contained $\mathcal{N}_{\mathrm{tot}} = 5000$ molecules at 300 K and $\rho = 0.8$. After 0.6 ns of equilibration, we ran a 0.2 ns production run and output

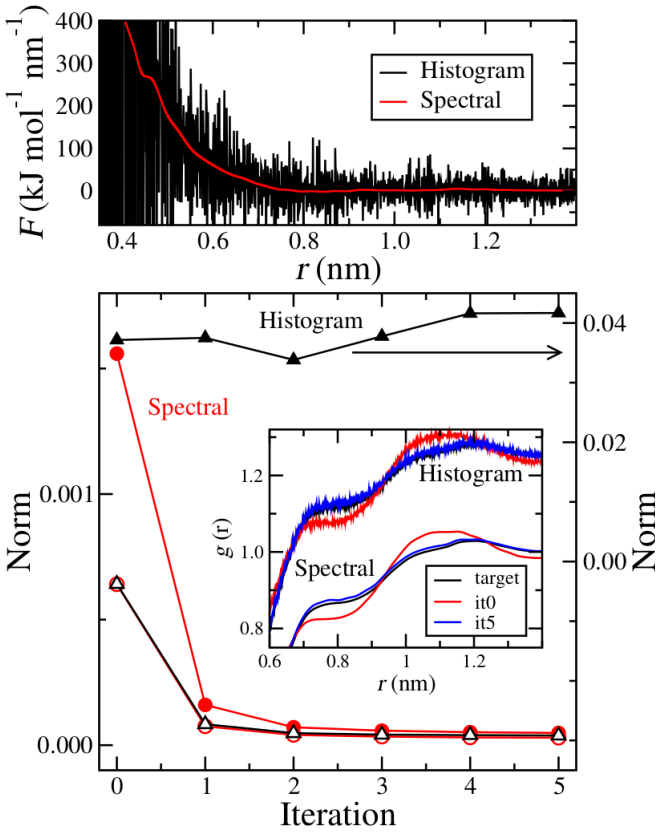

FIG. 2. *Top:* CG force for PS calculated via IBI with $n_c = 10^4$. The black curve (rough) is the histogram method result, whereas the red (smooth) curve is the SMC result. *Bottom:* $||g_i - g_t||^2_{L^2}$ (open symbols) and $||g_i - g_t||^2_{H^1}$ (closed symbols) norms for the histogram (circle) and SMC (triangle) methods as a function of IBI iteration. The inset shows the corresponding RDFs. Note that $||g - g_t||^2_{H^1}$ for the histogram method uses the right axis and is off the scale of the left axis.

configurations every 1 ps ($n_c = 200$). We take the corresponding 120-mode SMC reconstruction as a baseline for comparison, given its known convergence properties. For histogram-based approaches, we first partitioned the domain $0 \le r \le 0.9$ nm into 1800 intervals. After binning pair-separations from the first 20 frames, we used two separate smoothing algorithms to reduce noise: (i) a n-point moving mean with $n = 5$ and $n = 15$; and (ii) a Gaussian-kernel that convolves the histogram with $K(x) = \exp[-0.5(x/h)^2]$ for $h = 1$ pm and $h = 5$ pm. Figure 3 illustrates the key problem tied to the subjectivity of such methods: too little smoothing yields noisy RDFs (bottom inset), whereas too much washes out relevant features (top inset).

This figure also suggest that as a function of n_c, SMC converges to $g(r)$ more quickly than histogram-based approaches. To quantitatively test this, we estimated $g(r)$ for CG PS (cf. Figs. 1 and 2) as a function of n_c and computed the corresponding L^2 and and H^1 norms relative to the $n_c = 10^4$ case (which now acts as g_t). Figure 4 shows the results of this exercise. Most apparent, ev-

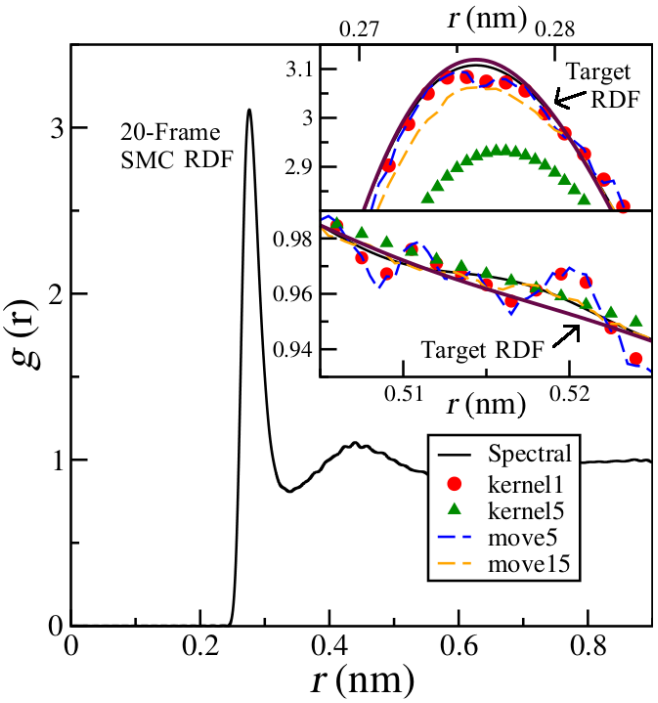

FIG. 3. Comparison of RDFs constructed using 20 frames of a 5000-molecule water simulation. The main figure shows that SMC captures both the sharp peak and the rapid transition around $r = 0.25$ nm. The insets compares the 20-frame SMC, kernel-smoothed, and moving-average RDFs relative to a 200-frame SMC RDF (dark purple). See main text for discussion.

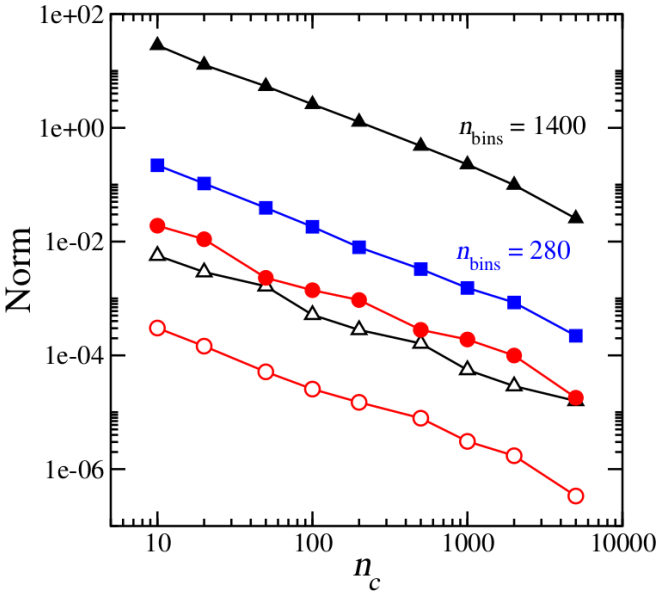

FIG. 4. $||g - g_t||^2_{L^2}$ and $||g - g_t||^2_{H^1}$ as a function of n_c for the 5th IBI update to the PS model in Fig. 2. Here g_t is the $n_c = 10^4$ RDF. Squares denote histogram estimates computed using 280 bins. All other symbols correspond to 1400 bins and have the same meanings as in previous figures.

ery norm decays as roughly $1/n_c$. Intuitively we expect this from the central limit theorem, since the variance in an average of N independent, identically distributed random variables should decay as the inverse of N. However, the SMC norms (circles) are at least an order of magnitude or more smaller than their histogram counterparts (squares and triangles). This suggests that the overhead required to generate pair separations can be reduced by a factor of 10 or more simply by using SMC.

Figure 4 also shows that increasing the histogram bin-width leads to *seemingly* smoother reconstructions of $g(r)$. This arises from the fact that more data points contribute to any given bin, thereby decreasing fluctuations. However, this *does not necessarily improve* the accuracy of such reconstructions, since bin counts are then averages taken over increasingly large domains.[5] Thus, the H^1 norm we propose should be used with caution, since it is likely not a valid assessment of histograms when the number of bins becomes too small. Along similar lines, we do not pursue quantitative comparison with convergence rates of smoothed histograms; such an analysis would require quantification of the uncertainties induced by smoothing, which can be highly non-trivial to estimate.

Analytically, the connection between SMC and histograms can be understood by framing the latter in the context of Eq. (5). Specifically, Eq. (5) reduces to a histogram bin count when the $\phi_j(r)$ are indicator functions $I_{[r_j, r_{j+1}]}$, i.e. constants on an interval (i.e. bin) $[r_j, r_{j+1}]$ and zero otherwise. These observations suggest that the ϕ_j act as a generalized histogram "bin." The fact that $\phi_j(r_k)$ may be non-zero for multiple j indicates that each pair separation r_k contributes to multiple "bins," albeit in unequal amounts.[6]

From a conceptual standpoint, binning is therefore equivalent to SMC insofar as Eq. (1) encompasses both approaches. Practically speaking, this suggests that both methods should be comparable in terms of computational time, which we generally find to be true. For the trajectories analyzed in this work, single-CPU binning computations take about 20 minutes or less using custom C++ codes, whereas their SMC counterparts take about three times as long for the same value of n_c. Given that MD simulations often take days, the real savings in our approach comes from needing orders of magnitude fewer pair-separations. Moreover, SMC is embarrassingly parallel, so that the relevant computations can be reduced to a matter of minutes on standard GPUs.

In concluding our discussion, we emphasize that de-

[5] In other words, increasing the bin width trades uncertainty along the vertical axis for uncertainty along the horizontal axis.

[6] That is, SMC bins data according to the characteristic wavelengths with which the r_k fall on the domain $[0, r_c]$.

spite its potential benefits, SMC nonetheless requires thoughtful implementation in order to be useful. In particular, spectral reconstructions may give slightly negative values of $g(r)$ for small separations near $r = 0$. This typically arises an incomplete destructive interference of the ϕ_j near the origin, but as Fig. 3 shows, such effects are often not visually apparent. Moreover, the problem is easily addressed by replacing Eq. (1) with an exponential decay $\exp(-r^k)$ near the origin (where k is a fit parameter), which yields a power-law force for small separations. We have found that this resolves any such issues with IBI when they arise.

Acknowledgments: the authors thank Timothy Burns, Andrew Dienstfrey, and Vincent Shen for useful feedback during preparation of this manuscript. *This work is a contribution of the National Institute of Standards and Technology and is not subject to copyright in the United States.*

[1] J. G. Kirkwood and F. P. Buff, The Journal of Chemical Physics **19** (1951).
[2] K. E. Newman, Chem. Soc. Rev. **23**, 31 (1994).
[3] M. Allen and D. Tildesley, *Computer simulation of liquids*, Oxford science publications (Clarendon Press, 1987).
[4] J. Yarnell, M. Katz, R. Wenzel, and S. Koenig, Physical Review A **7**, 2130 (1973).
[5] N. Ashcroft and N. Mermin, *Solid State Physics*, HRW international editions (Holt, Rinehart and Winston, 1976).
[6] C.-C. Fu, P. M. Kulkarni, M. Scott Shell, and L. Gary Leal, The Journal of Chemical Physics **137**, 164106 (2012), http://dx.doi.org/10.1063/1.4759463.
[7] C. Peter and K. Kremer, Soft Matter **5**, 4357 (2009).
[8] B. Bayramoglu and R. Faller, Macromolecules **45**, 9205 (2012).
[9] F. Muller-Plathe, ChemPhysChem **3**, 754 (2002).
[10] W. G. Noid, Journal of Chemical Physics **139** (2013).
[11] D. Reith, M. Putz, and F. Muller-Plathe, Journal of Computational Chemistry **24**, 1624 (2003).
[12] B. G. Levine, J. E. Stone, and A. Kohlmeyer, J. Comput. Phys. **230**, 3556 (2011).
[13] D. Frenkel and B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications*, Computational science series (Elsevier Science, 2001).
[14] L. Evans, *Partial Differential Equations*, Graduate studies in mathematics (American Mathematical Society, 2010).
[15] P. Biswas, D. N. Tafen, F. Inam, B. Cai, and D. A. Drabold, Journal of Physics: Condensed Matter **21**, 084207 (2009).
[16] H. Khandelia, A. A. Langham, and Y. N. Kaznessis, Biochimica et Biophysica Acta (BBA) - Biomembranes **1758**, 1224 (2006).
[17] C. P. Robert and G. Casella, "Introducing monte carlo methods with r," (Springer New York, New York, NY, 2010) Chap. Monte Carlo Integration, pp. 61–88.
[18] L. Lin, Y. Saad, and C. Yang, SIAM Review **58**, 34 (2016).
[19] J. Boyd, *Chebyshev and Fourier Spectral Methods: Second Revised Edition*, Dover Books on Mathematics (Dover Publications, 2001).
[20] W. Strauss, *Partial Differential Equations: An Introduction* (Wiley, 1992).
[21] H. Sun, S. Mumby, J. Maple, and A. Hagler, Journal of the American Chemical Society **116**, 2978 (1994).
[22] T. C. Moore, C. R. Iacovella, and C. McCabe, Journal of Chemical Physics **140** (2014).
[23] R. Faller, Polymer **45**, 3869 (2004).
[24] W. Jorgensen, J. Chandrasekhar, J. Madura, R. Impey, and M. Klein, Journal of Chemical Physics **79**, 926 (1983).